Foraging for Trust: Exploring Rationality and the Stag Hunt Game*

Steven O. Kimbrough
565 JMHH
kimbrough@wharton.upenn.edu
215-898-5133

iTrust Workshop, 25 May 2005

These foils available online at `http://opim-sun.wharton.upenn.edu/~sok/sokpapers/2005/itrust/trust-foils.pdf`

---

[*] File: trust-foils.tex/pdf.

Cooperation: The Cement of Society[?]

- Cooperation: key to social felicity (whether a good thing or not)
- Trust: a means to achieve cooperation. Assumption of risk, depending on behavior of counter-parties. One trusts that the other guy will do the right thing.
- Is it like altruism, perhaps: more apparent than real?
- Must ask: Cooperation or trust for what?
  - Usually: for mutual advantage. Pareto optimality.
  - Other or additional goals? Fairness? Justice? Team purposes? Agreed goals? Game theory has always abstracted these away.

Institutions as Substitutes for Trust

- Insurance
- Contracts
- Social structure (not usually counted as an institution)
- Limitations of the above
- Note broader context of political economy: Debate between institutionalism (‘Hobbes’), and self-organization and laissez-faire (’Smith’).

Puzzles Regarding Trust

- What sustains it?
- What creates it?
- What destroys it?
- What are its dynamics?

Even broader scientific context: “Birds do it, bees do it, even monkeys up in trees do it.” And we do it. My frame of reference here.

Larger Puzzles: Game Theory

[G]ame theory does not pretend to tell you how to make judgments about the shortcomings of an opponent. In making such judgments, you would be better advised to consult a psychologist than a game theorist.

Game theory is about what players will do when it is understood that both are rational in some sense. Sometimes, . . . this means that an orthodox game-theoretic analysis is not necessarily a very helpful guide on how to play against real people. . . . It is however true that, unless there are good reasons for supposing that the people involved will behave rationally, game theory cannot realistically be used in a naive way to make predictions about what real people will do. As a consequence, a player would often be unwise to use the strategy that a game theory book may label as “optimal” because this will usually only optimal if *everyone* plans to play optimally.

Of course, there are circumstances in which is *is* reasonable to work on the hypothesis that people will behave in a reasonably rational manner. Economics is somewhat shakily founded on the assumption that this will typically be the case in commercial and business transactions. However, it would be skating on very thin ice to use game theory for predictive purposes if none of the following criteria were satisfied:

- The game is simple.

- The players have played the game many times before and hence have had much opportunity for trial-and-error learning.

- The incentives for playing well are adequate.

—*Fun and Games*, pp. 50–1, Ken Binmore, 1992.

Trouble in River City: Game Theory

Classical game theory focuses on the equilibria of games (with the Nash equilibrium being a fundamental, organizing concept), and makes very strong assumptions in investigating an ideal form of rationality.

> Equilibrium analysis [i.e., for our purposes classical game theory] is based formally on the presumptions that every player maximizes perfectly and completely against the strategies of his opponents, that the character of those opponents and their strategies are perfectly known, and that players are able to evaluate all their options. [Kre90b, page 139]

Kreps adds immediately that "None of these conditions will be met in all respects in reality. The behaviour of individuals in economic contexts may approximate these assumptions in some contexts, but the approximation may be insuffient in others.”

See [Kre90a] for elaboration.

Trouble (con't.)

Ironically, [classical] game theory is often hoisted on its own pétard: many of its most fundamental predictions—predictions that would have been too vague to test with any confidence in the pre-game-theoretic era—are *decisively and repeatedly disconfirmed,* in laboratory settings, with substantial agreement among experimenters, regardless of their theoretical priors. [Gin00, page xxiv]

Copernican revolution

Evolutionary game theory effects something of a Copernican revolution in the study of games. The change is captured nicely in the following passage from a recent textbook:

> . . . *game theory is about the emergence, transformation, diffusion, and stabilization of forms of behavior.* Traditionally, game theory has been seen as a theory of how "rational agents" *do* behave, and/or how the rest of us *should* behave. Ironically, game theory which for so long was predicated upon agent rationality, has shown us, by example, the shakiness of the concept. For one thing, the centipede game and others like it show that there is nothing substantively "rational" about even so simple a thing as eliminating dominated strategies

. . . . Moreover, the solution to some games (even when unique) is often so sophisticated that it is implausible that ordinary people would be willing to spend the resources to discover it. This supports the evolutionary notion that good strategies diffuse across populations of players rather than being learned by "rational optimizers." Finally, experimental studies of dictator, ultimatum, and public goods games indicate that if people are "rational," it must be in a sense far more sophisticated than the simple, self-interested, maximization of expected utility.

It is better to drop the term "rational" altogether, which is what we do in this book . . . .

In short, evolutionary game theory replaces the idea that games have "solutions" that agents "learn," with the idea that games are embedded in natural and social processes that produce agents who play effectively.

> Dispensing with the rationality postulate does not imply that people are *irrational* (whatever that means). The point is that the concept of "rationality" does not help us understand the world. [Gin00, page xxv-xxvi]

- I'm going to take a different tack and suggest that a concept of rationality is indeed very useful, but that this has to be a different sort of rationality than that assumed in classical game theory.
- I'll develop these ideas in the context of trust and the Stag Hunt game.

Fundamental Problem: The Folk Theorem

- Genuine theorem(s).
- Roughly: in a *repeated game* nearly any outcome can be supported by a Nash equilibrium.
- Repeated game: indefinitely repeated play among the same players.

  Affords opportunities to signal, learn, and collude.
- Iterated game: multiple plays of a game, but with different players.
- Thus, in repeated play, the Nash equilibrium becomes a trivial prediction and hence uninteresting.

Models for Trust

- Actions that do-or-do-not involve trust are essentially strategic, and so addressable by the theory of games. (Decisions: parametric versus strategic.)
- On the strategy of modeling (in this context)
- What are the games?
- KISS: Prisoner's Dilemma and Stag Hunt

Prisoner's Dilemma

	Cooperate (C)	Defect (D)
Cooperate (C)	3 3	4 0
Defect (D)	0 4	1 1

Required: $T > R > P > S$ and $2R > T + S$.

	Cooperate (C)	Defect (D)
Cooperate (C)	R' R	T' S
Defect (D)	S' T	P' P

Stag Hunt

	Hunt stag (S)	Chase hare (H)
Hunt stag (S)	4 4	3 0
Chase hare (H)	0 3	1 1

Generally: $R > T \geq P > S$. Three Nash equilibria in the one-shot game.

	Hunt stag (S)	Chase hare (H)
Hunt stag (S)	R′ R	T′ S
Chase hare (H)	S′ T	P′ P

Stag Hunt story

> The French philosopher, Jean Jacques Rousseau, presented the following situation. Two hunters can either jointly hunt a stag (an adult deer and rather large meal) or individually hunt a rabbit (tasty, but substantially less filling). Hunting stags is quite challenging and requires mutual cooperation. If either hunts a stag alone, the chance of success is minimal. Hunting stags is most beneficial for society but requires a lot of trust among its members. [Gam05]

Presents a delicate dilemma: Play it safe or take a risk with the prospect of a higher reward? Should you, will you, trust?

Generally in one-shot games played iteratively, subjects learn not to trust. Coheres with existing theory. (Depends, however, on payoffs [BSV00].)

Other modeling approaches: Repeated play of Stag Hunt

- First, focus on representative results, reported by Brian Skyrms, [Sky01, Sky04].
- Then, results in my paper in the proceedings.
- Then, new results, since the paper
- Then, conclude with remarks about rationality and the research programme.

Replicator Dynamics

- A basic, simple evolutionary model, widely explored. Random ("mean field") play between pairs of individuals (=strategies) in a population. Points are totaled and a strategy's chance of being present in the next generation is proportional to its success in the present generation.

- Let the two possible strategies be S (hunt stag) and H (hunt hare). Seed a population and let them go. Broadly, you get one or the other strategy going to fixation, usually H. 50:50 start usual leads to all H [Sky04, chapter 1]. Each is an ESS. Mutation can bounce you out, but it will typically take a very long time.

  Generally: disappointing findings (cf. also [KMR93, HDSH01]).

Added analytic note

Assuming a symmetric game:

	x: Hunt stag (S)	$(1-x)$: Chase hare (H)
Hunt stag (S)	R R	T S
Chase hare (H)	S T	P P

The S is favored iff

$$Rx + S(1-x) > Tx + P(1-x) \tag{1}$$

@ R=4, T=3, P=1, S=0, $x = \frac{1}{2}$.

Playing with Neighbors

- Play in a circle leads to results not much different than the replicator dynamics [Ell93]. And it happens quickly. Based on *best-response* update policy.

- Skyrms and Pemantle [SP00], however, find considerable levels of trust achieved in Stag Hunt, if the agents can exercise some choice in whether or not to play with a given counter-party.

  This is a (social) network model, in which agents are nodes and playing relationships are arcs. The S agents learn to get the other S agents into their social networks.

2D Models: The lattice or gridscape

6×6 gridscape:

	1	2	3	4	5	6
1						
2		NW	N	NE		
3		W	X	E		
4		SW	S	SE		
5						
6						

X plays eight neighbors (Moore neighbohood). Update rule: *imitate-the-best.* As Skyrms reports [Sky04, chapter 3], S typically takes over.

Materials and results from the conference paper

	1	2	3	4	5	6	7	8
1								
2		S	S	S				
3		S	S	S				
4		S	S	S				
5								
6								
7								

(a)
Generation
x

	1	2	3	4	5	6	7	8
1			S					
2		S	S	S				
3	S	S	S	S	S			
4		S	S	S				
5			S					
6								
7								

(b)
Generation
x+1

Continuing

	1	2	3	4	5	6	7	8
1		S	S	S				
2	S	S	S	S	S			
3	S	S	S	S	S			
4	S	S	S	S	S			
5		S	S	S				
6								
7								

(c) Generation x+2

Analytically

- Given imitate-the-best, then
- if we get a 3×3 or larger block of Ss, and
- if $5(R-P) > 3(T-S)$ (true when R=4, T=3, P=1, S=0), then
- the chance of takeover by S is quite high. (See paper for details.)

In sum on spatial (network) models for Stag Hunt

- There are known variations and much to be investigated, but . . .
- At least some models show that under reasonable conditions trust in Stag Hunt can spontaneously arise and be maintained.
- Key factors are:

  (a) payoff/reward structure — which is neglected in the classical theory

  (b) social structure — also classically neglected
- Results for Prisoner's Dilemma, too.

Now, models of individual learning

- Replicator dynamics (and most other mainstream models): adaptation, or evolution, but not learning.
- Network models: imitation learning; limited but not to be neglected. Richer models have begun to be explored.

  Upcoming, LPS models: generalize and abstract earlier models:
- Reinforcement learning
- MLPS, a Markov model
- Models of matching and other adaptive learning processes

Utility theory; rational choice theory

Begins by assuming a set Ω of outcomes and a preference relation $\succeq$ on Ω.

Posits for all $a, b, c \in \Omega$

1. Totality: $a \succeq b$ or $b \succeq a$.
2. Transitivity: if $a \succeq b$ and $b \succeq c$, then $a \succeq c$.

Sounds good. Why not?

Foraging context

n patches each of which may be tried for a reward. At the next time step the agent may stay or move to another patch. NB: multi-armed bandits.

Problem: where does $\succeq$ come from?

It is determined by context-specific conditions. The agent's challenge is less to reason about $\succeq$ than to discover it, to ascertain it with sufficient precision to support effective action.

$\Rightarrow$ Utility theory, rational choice theory, is a tool of minimal use in a foraging context.

Rational choice theory is appropriate for choosing from Ω given a well-articulated and supported $\succeq$.

Exploring rationality

- A *theory of exploring rationality* is needed and appropriate for choice when $\succeq$ is poorly known.

- What would that look like?

  Compare maxims of rationality:

  1. "It is rational to have transitive preferences."
  2. "In the presence of incomplete knowledge, is it rational to probe one's environment, sampling for useful information."

  Each maxim commands consideration, but in distinct contexts.

- Illustration: a multi-armed bandit problem in which a casino customer must sample the slot machines for the best chance of net positive payoff.

Contrasting summary

- EU or Rational Choice theory is a theory of *maximum taking* (MT) among presented and well-evaluated options.
- A theory of exploring rationality is a theory of *maximum seeking* (MS) among incompletely known options.

  Inherently a trail-and-error process.
- Rational Choice theory and exploring rationality theories address different problems.

Q: Why do you hate Hollywood? A: [Robert Altman] I don't hate Hollywood. They make shoes. I make gloves. We're in different businesses.

Rejoinder: Bayes & SEU

You forget *subjective* EU theory. This with, Bayesian revision, solves the problem. It is a universal approach to decision making.

Rejoinder rejoinder:

- Here we have to be brief, but I think not; the rejoinder doesn't work. There are multiple problems.
- Bayesian revision is often impractical; too difficult or costly to undertake, especially for limited beings. We need to investigate simple, robust heuristics.
- It is far from clear that SEU is universal. In particular, finding a

nonparametric prior distribution is quite problematic. Inevitably, arbitrary assumptions are made.

- SEU fails to deliver what we (usually) need: an objectively grounded prescription for action. Why my beliefs? Why not yours? Or Fred's cat? Or readings of tea leaves?

Pragmatism with instrumentalism is a philosophy closer to the mark. Beliefs and concepts are to be seen primarily as tools for negotiating the world, and should be assessed as such.

I'm saying that S/EU theory is not an appropriate, not a very useful tool for foraging contexts. For these we need an exploring rationality.

The trick is to find and assess simple, robust models for what are, broadly, foraging contexts. Now we'll apply this to games, and to the Stag Hunt in particular. First, a *class* of such models: LPS models.

LPS: Learning in Policy Space, pseudo code

repeat forever:

1. Select a policy $\pi_i \in \Pi$, where Π is the consideration set of policies.
2. Pick a length of play, l, for policy π_i.
3. Play the next l rounds of the game using π_i.

   Note: At each round, π_i will observe the current state, s_t, take an action a and obtain a reward r_t.
4. Update V^{π_i} based on the individual-round rewards, r_ts, obtained during the l rounds of play of policy π_i. Loop.

Stag Hunt with reinforcement (Q-) learning

See [KL05, KLK04].

	C_L (Stag)	C_R (Hare)
R_U (Stag)	(5,5)*#	(0,3)
R_D (Hare)	(3,0)	(δ, δ)#

Figure 1: Parametric Stag Hunt

Representative results

Table 1: Summary of results for Stag Hunt. ε-greedy action selection. Totals for the last 100 rounds of 100 series of 10,000 plays.

SS	SH	HS	HH	δ	Row's % CC
9390	126	122	362	0	0.978
9546	91	108	255	0.5	0.976
9211	112	125	552	0.75	0.975
8864	119	110	907	1	0.975
8634	115	132	1119	1.25	0.971
7914	122	130	1834	1.5	0.963
7822	122	104	1952	2	0.965
5936	87	101	3876	2.5	0.925
5266	121	106	4507	3	0.736

MLPS: A model for exploring rationality

- MLPS: Markov learning in policy space. See the conference paper for details.
- Basic setup:
  - The supergame: individual rounds of play of Stag Hunt repeated indefinitely, between two players.
  - The supergame is divided into games: rounds of play of Stag Hunt, each of fixed length.
  - The games are divided into epochs: rounds of play of length le.
    Each game consists of n_e epochs.

Individual play in the MLPS model

- Each player has a *consideration set of policies for play,* $\mathcal{S}$.

  E.g., Always hunt stag, Always hunt hare, Tit for Tat, Suspicious Tit for Tat

- At the beginning of each game, each player independently picks a *focal strategy* from its $\mathcal{S}$. Choice is made by fitness-proportional selection on returns from the last game for policies in $\mathcal{S}$.

- At the beginning of each epoch in a game, each player independently picks a *policy-in-use* for the duration of the epoch. With probability $1-\varepsilon$, a player picks its current focal policy. With probability ε the player randomly chooses among the non-focal policies.

Example (from the paper)

- Each player has two policies in its $\mathcal{S}$: Always Defect and Tit for Tat. $l_e = 10$, $\varepsilon = 0.1$.

- Assume: n_e large enough, or mechanism-based coordination strong enough that expected values of returns to policies are exactly realized.

- We have then a Markov process, with state transitions occurring a the end of games. There are 4 states: (1) Row focuses on Tit for Tat & Colum focuses on Tit for Tat, (2) Row focuses on Tit for Tat & Colum focuses on Always Defect, . . .

  Note: Neither player knows what state the system is in.

The transition matrix

	s(1)=(1,1)	s(2)=(1,2)	s(3)=(2,1)	s(4)=(2,2)
s(1)	$0.7577 \cdot 0.7577$ $= 0.5741$	$0.7577 \cdot 0.2423$ $= 0.1836$	$0.2423 \cdot 0.7577$ $= 0.1836$	$0.2423 \cdot 0.2423$ $= 0.0587$
s(2)	$0.5426 \cdot 0.7577$ $= 0.4111$	$0.5426 \cdot 0.2423$ $= 0.1315$	$0.4574 \cdot 0.7577$ $= 0.3466$	$0.4574 \cdot 0.2423$ $= 0.1108$
s(3)	$0.7577 \cdot 0.5426$ $= 0.4111$	$0.7577 \cdot 0.4574$ $= 0.3466$	$0.2423 \cdot 0.5426$ $= 0.1315$	$0.2423 \cdot 0.4574$ $= 0.1108$
s(4)	$0.5426 \cdot 0.5426$ $= 0.2944$	$0.5426 \cdot 0.4574$ $= 0.2482$	$0.4574 \cdot 0.5426$ $= 0.2482$	$0.4574 \cdot 0.4574$ $= 0.2092$

Table 5: Stag Hunt transition matrix data assuming fitness proportional policy selection by both players. Numeric example for $\varepsilon = 0.1 = \varepsilon_1 = \varepsilon_2$.

At convergence of the Markov process

Pr(s(1))	Pr(s(2))	Pr(s(3))	Pr(s(4))
0.4779	0.2134	0.2134	0.0953

So 90%+ of the time at least one agent is playing $\mathrm{T\scriptsize{F}\normalsize{T}}$. They learn to trust (a lot). Note the expected take for Row per epoch by state:

1. $(1-\varepsilon)(40-31\varepsilon)+\varepsilon(12-2\varepsilon)=34.39$
2. $(1-\varepsilon)(9+31\varepsilon)+\varepsilon(10+2\varepsilon)=11.91$
3. $\varepsilon(40-31\varepsilon)+(1-\varepsilon)(12-2\varepsilon)=14.31$
4. $\varepsilon(9+31\varepsilon)+(1-\varepsilon)(10+2\varepsilon)=10.39$

The players do better playing this way

- Row (and Column) gets 2.302 per round of play (on average).
- At the (H,H) Nash equilibrium, each player gets 1 per round of play.
- If both players play AllD with ε-greedy exploration to Tit for Tat, each gets 1.039 per round of play.
- If both players play the mixed Nash equilibrium, each gets on average 2 per round of play.

An intuitive, very simple, cognitively undemanding model yields substantial trust and cooperation as an emergent property. Compare with 'rational fools'.

How robust and realistic is the MLPS model?

- Very robust, across a variety of settings and games. See paper and working paper cited.

- In detail, not realistic at all, but
  - It captures some important features (policy space, trial and error exploring)
  - Speedy convergence
  - Simulation with relaxed assumptions leads to qualitatively similar results.

- It demonstrates that simple learning procedures of a certain type can reliably produce a degree of trust and cooperation.

Another model: Matching law and variants

- See [Kim05], "A Note on the Matching Law and Repeated Games in Strategic Form."
- The *matching law*: highly robust and confirmed law in psychology, championed by Richard Herrnstein [Her97], models learning behavior. (Yes, *that* Richard Herrnstein.)

Let the B_is be the behavior alternatives, with the R_is as the associated reinforcements. The matching law says that at convergence of learning:

$$\frac{B_i}{B_1 + B_2 + \ldots + B_n} = \frac{R_i}{R_1 + R_2 + \ldots + R_n} \tag{2}$$

Getting there

The associated hypothesis of *melioration* [Her97] addresses the dynamics of learning under the matching law. According to this hypothesis [Her97, page 77], the adapting/learning agent adjusts its B_is so that:

$$\frac{R_1}{B_1} = \frac{R_2}{B_2} = \ldots = \frac{R_n}{B_n} \tag{3}$$

It is assumed throughout that the sum of the B_is is fixed, at least relatively so. In the case of two alternatives, doing more of B_1 entails doing less of B_2, for example.

Can the matching law contribute to an explanation of trust?

- Developed for parametric, not strategic decisions. Let's look at its behavior in a game, in Stag Hunt.

	y: Hunt Stag (S)	$(1-y)$: Hunt Hare (H)
x: Hunt Stag (S)	4,4	0,2
$(1-x)$: Hunt Hare (H)	2,0	1,1

Figure 2: Example Stag Hunt game

Let x be the probability that Row hunts stag (and $(1-x)$ the probability Row hunts hare), and y similarly for Column. In this particular game, there are three Nash equilibria: $(x=1,y=1),(x=0,y=0)$ and the mixed equilibrium $(x=1/3,y=1/3)$.

Can agents learn (this way) to trust?

What will the players learn to do when playing each other iteratively? Since the game is symmetric, we can limit our attention to Row's perspective. From that perspective, the expected rate of return for hunting stag (S) is $4y$ and the expected rate of return for hunting hare (H) is $1 + y$. At the Nash equilibrium in mixed strategies, of course, the expected rates of return are equal: at $y = 1/3$, $4y = 1 + y$. Off equilibrium, i.e., for $y \neq 1/3$, Row will find one or the other behaviors—S or H—comparatively more attractive. If $y > 1/3$, Row will find the return rate from S higher than the return rate from H. Similarly, if $y < 1/3$ Row will find H more attractive than S. Given a fixed value of y (here, the case is parametric, not strategic), the matching law predicts that Row will learn to play S always if $y > 1/3$ and H always if $y < 1/3$.

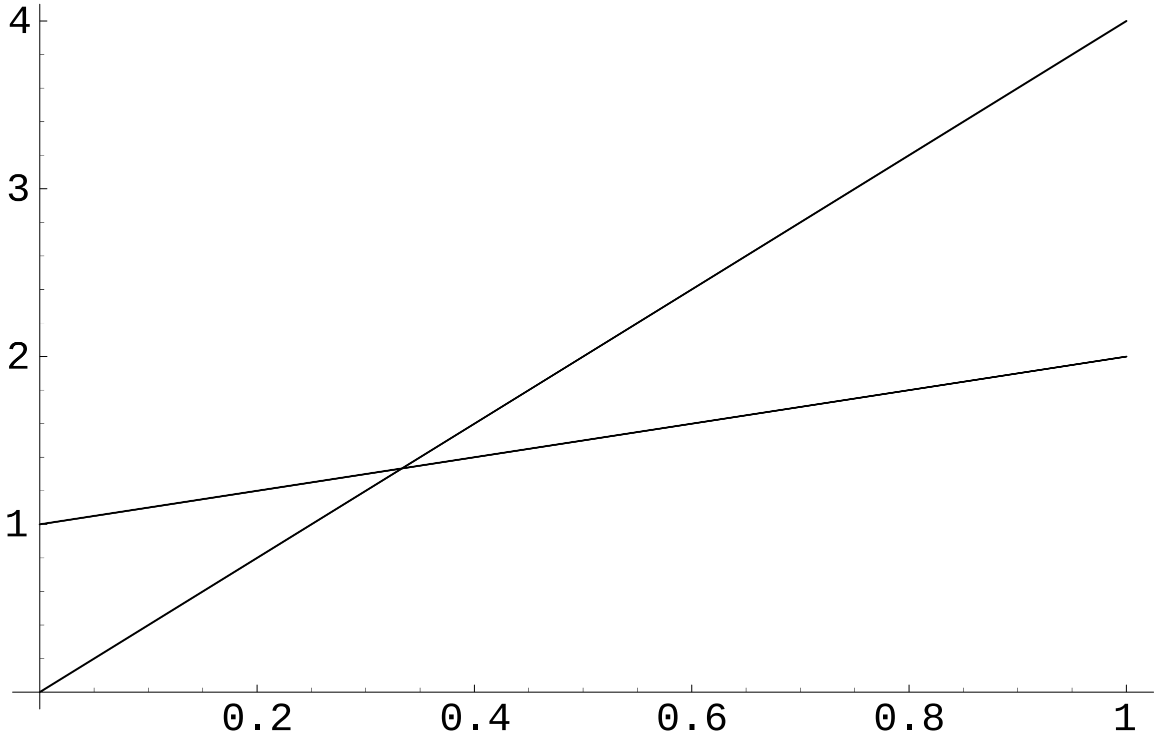

Figure 3: `rewards = Plot[{4y, (y + 1)}, {y, 0, 1}]`, `Display["rewards.eps", rewards, "EPS"]`

So, melioration and matching may or may not lead to trust

- They will if both players start playing S with probability $> \frac{1}{3}$.
- They will not if both players start playing S with probability $< \frac{1}{3}$.
- Is there some other simple, plausible, and realistic rule or policy that might help?

Fitness proportional selection

Fitness proportional sampling of behaviors is a simple and intuitively natural policy. Under this policy, the probability of engaging in behavior B_i as the next behavior is just the ratio of R_i, the currently estimated reward for B_i to the sum of the R_is. In expectation for our Stag Hunt game, then

$$\text{Pr}(S) = \frac{4y}{(4y+1+y)} = \frac{4y}{(5y+1)} \tag{4}$$

$$\text{Pr}(H) = \frac{(1+y)}{(5y+1)} \tag{5}$$

If both Row and Column are adjusting x and y by fitness proportional sampling, the system will be driven to $x = y = 3/5$. At this convergence point Row gets $52/25 = 2.08$, as does Column.

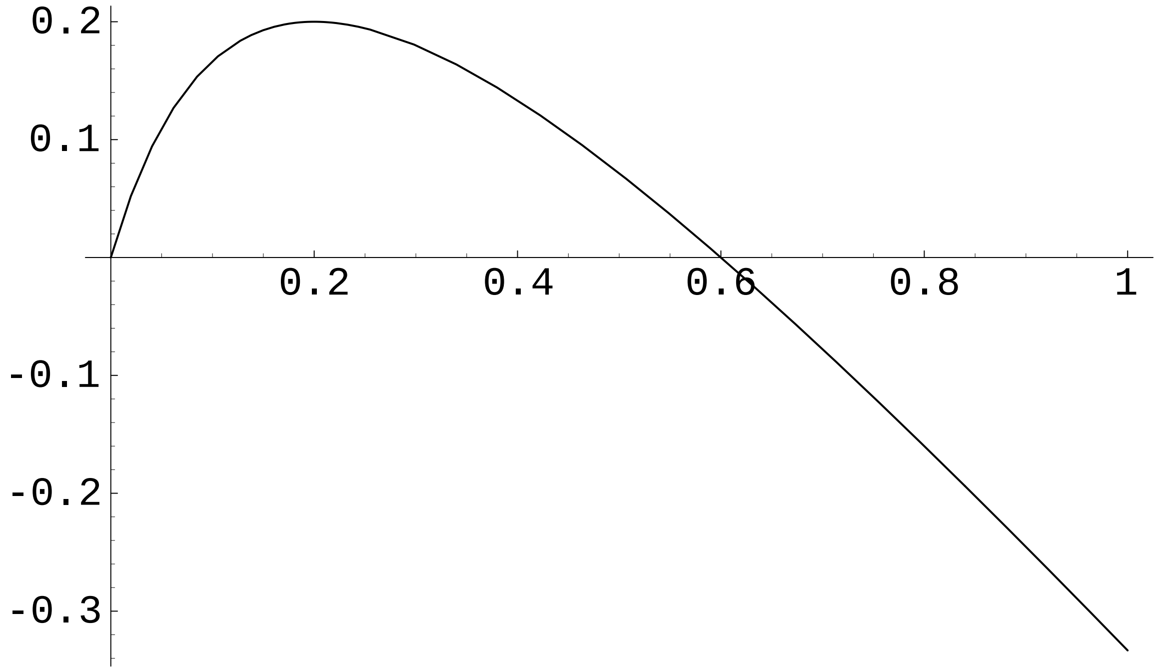

Figure 4: `differenceplot = Plot[(4y/(5y + 1)) - y, {y, 0, 1}]`, `Display["differenceplot.eps", differenceplot, "EPS"]`

Combining scenario

- At the start of learning, both agents have very little knowledge of the situation and so engage in fitness proportional selection of play (S or H).
- This drives them towards $x = y = 0.6$ and an equilibrium.
- Noticing the equilibrium, the players switch to matching and melioration, and are driven thus to $x = y = 1$. They learn to trust.

Smarter play can pay off

- Suppose the players divide play into epochs of length 10, and consider TIT FOR TAT (instead of ALWAYS HUNT STAG) and as before ALWAYS HUNT HARE.

  The new stage game—still a Stag Hunt—is

	TIT FOR TAT (S)	ALLD (H)
TIT FOR TAT (S)	40,40	9,11
ALLD (H)	11,9	10,10

Figure 5: Stag Hunt game payoffs per epoch of length 10

New equilibria

- As always, 3 in the one-shot Stag Hunt. Here they are at: $x = y = 1$, $x = y = 0$, and $x = y = \frac{1}{30}$
- An odd, even paradoxical, fact about the mixed Nash equilibrium is that as S becomes more attractive to both players in the payoff matrix, it becomes *less* frequently visited in the mixed Nash equilibrium.
- Note, however, that the new region in which matching drives the players to cooperation is now much larger. Formerly: $x, y > \frac{1}{3}$. Now: $x, y > \frac{1}{30}$.
- Further, fitness proportional selection does better, too. See Figure 6. The equilibrium has moved from 3/5 to 3/4.

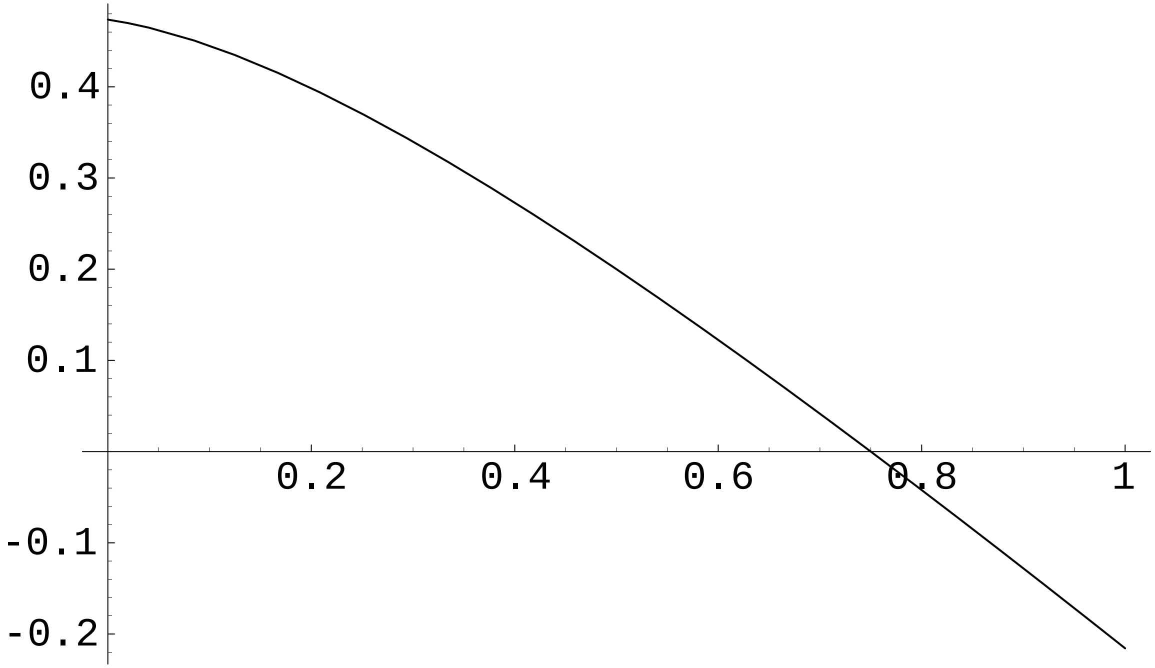

Figure 6: `differenceplotTfT = Plot[((40c + (1 - c) 9)/((40c + (1 - c)9) + (11c + (1 - c)10))) - c, {c, 0, 1}], Display["differenceplotTfT.eps", differenceplotTfT, "EPS"]`

Summary & Discussion

- How can/does trust arise and what sustains/destroys it?

  Huge and hugely important questions. Only a few small contributions here. Overview of a certain perspective and tabling of many research questions.

- Perspective: fundamentally strategic as distinct from a parametric decisions by agents. Game theory.

- But classical game theory is problematic here in many ways, discussed above.

Summary & Discussion

- Aside.

  Not discussed above: classical theory not sensitive to non-ordinal aspects of payoff sizes. Counter-intuitive. The models here discussed are sensitive in this way and quite plausibly so.

- Rationality, as conceived by EU (rational choice) theory, turns out to be a problematic concept.

  Evolutionary game theory eschews it.

  I want to argue that it is useful to think in terms of a different sort of rationality, an *exploring rationality*, in these contexts.

Summary & Discussion

- LPS models exemplify exploring rationality.
- We've looked at several such models in the context of the Stag Hung game: reinforcement learning, MLPS, matching, fitness proportional selection, and combinations of the last two.
- In each case, trust emerges and is sustained robustly.

  Not so either in the classical theory or in the (mean field) replicator dynamics.
- The LPS models have in addition a number of attractive features.

  Simple. Computable. Sensitive to the payoff structure, and plausibly.

Larger picture(s)

“The key to the evolution of cooperation, collective action, and social structure is correlation. Correlation of interactions allows the evolution of cooperative social structure that would otherwise be impossible. Social institutions enable, and to a large part consist of, correlated interactions.”

—*Stag Hunt and the Evolution of Social Structure*, Skyrms, [Sky04, xii]

To which we must add:

- Individual learning regimes (the LPS family), and
- Sophistication in policies of play (the consideration set)

as key factors in explaining trust, cooperation, and the social order, generally.

Ahead

- This is progress, but very much remains to be done if we are to transition from models that proffer ‘insight’ to models that may be used for, e.g., policy making.

- Perhaps next: models of this sort for particular types of markets, with a small number of players and substantial repititon, e.g., spot markets for electricity,

- Eventually, models in which agents operate in an “ecology of games” [Lon58], gleaning experience constantly, and using analogy and similarity mapping to deal sensibly with new situations.

References

[BSV00] Raymond Battalio, Larry Samuelson, and John Van Huyck, *Optimization incentives and coordination failure in laboratory stag hunt games*, WWW: `http://erl.tamu.edu/JVH_gtee/C4.htm`, May 2000, Accessed 2005/05/09.

[Ell93] Glenn Ellison, *Learning, local interaction and coordination*, Econometrica **61** (1993), 1047–1071.

[Gam05] GameTheory.net, *Stag hunt*, `http://www.gametheory.net/Dictionary/Games/StagHunt.html`, Accessed 8 February 2005.

[Gin00] Herbert Gintis, *Game theory evolving: A problem-centered introduction to modeling strategic interaction*, Princeton University Press, Princeton, NJ, 2000.

[HDSH01] Christoph Hauert, Silvia De Monte, Karl Sigmund, and Josef Hofbauer, *Oscillations in optional public good games*, `http://www.iiasa.ac.at/Publications/Documents/IR-01-036.pdf`, September 2001, Accessed 2005/05/09.

[Her97] Richard J. Herrnstein, *The matching law: Papers in psychology and economics*, Harvard University Press, Cambridge, MA, 1997.

[Kim05] Steven O. Kimbrough, *A note on the matching law and repeated games in strategic form*, working paper, posted on the Web, April 2005, `http://opim-sun.wharton.upenn.edu/~sok/sokpapers/2006/note-matching-law.pdf`.

[KL05] Steven O. Kimbrough and Ming Lu, *Simple reinforcement learning agents: Pareto beats Nash in an algorithmic game theory study*, Information Systems and e-Business **3** (2005), no. 1, 1–19.

[KLK04] Steven O. Kimbrough, Ming Lu, and Ann Kuo, *A note on strategic learning in policy space*, Formal Modelling in Electronic Commerce: Representation, Inference, and Strategic Interaction (Steven O. Kimbrough and D. J. Wu, eds.), Springer, Berlin, Germany, 2004, pp. 463–475.

[KMR93] M. Kandori, G. Mailath, and R. Rob, *Learning, mutation, and long run equilibria in games*, Econometrica **61** (1993), 29–56.

[Kre90a] David M. Kreps, *A course in microeconomic theory*, Princeton University Press, Princeton, NJ, 1990.

[Kre90b] ———, *Game theory and economic modeling*, Clarendon Press, Oxford, England, 1990.

[Lon58] Norton E. Long, *The local community as an ecology of games*, American Journal of Sociology **LXIV** (1958), no. 3, 251–61.

[Sky01] Brian Skyrms, *The stag hunt*, World Wide Web, 2001, Proceedings and Addresses of the American Philosophical Association. `http://www.lps.uci.edu/home/fac-staff/faculty/skyrms/StagHunt.pdf`. Accessed September 2004.

[Sky04] ———, *The stag hunt and the evolution of social structure*, Cambridge University Press, Cambridge, UK, 2004.

[SP00] Brian Skyrms and Robin Pemantle, *A dynamic model of social network formation*, Proceedings of the National Academy of Sciences **97** (August 1, 2000), no. 16, 9340–9346.

`$Id: trust-foils.tex,v 1.6 2005/05/14 00:51:37 sok Exp $`